Reduced-Complexity Models for Network Performance Prediction

R. Srikant and Bruce Hajek
University of Illinois

Models

- Stochastic Queueing Models at the Packet Level
- Fluid Models at the Flow Level
- Loss Models at the Connection Level
- Game-Theoretic Models of a Federation of Networks

Recent Results

- AVQ: An adaptive virtual queue algorithm for active queue management (Kunniyur, S)
- Max-Min Fair Multicast Congestion Control (Graves, S, Towsley) (Workshop I)
- Fair Multicast Congestion Control for General Fairness Criteria (Deb, S)
- Comments on Pricing/Performance of Federated Networks (Gopalakrishnan, H)

I. Adaptive Virtual Queue

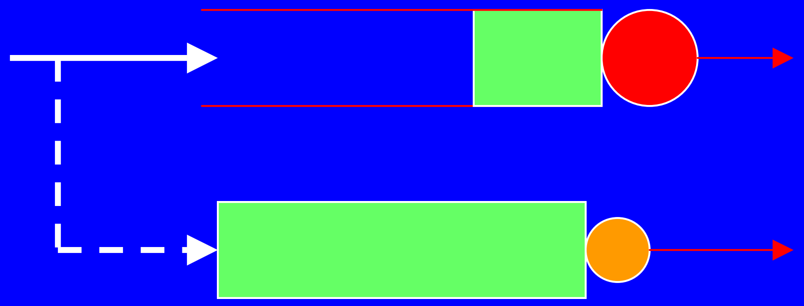

- Mark packets when virtual queue is full
- How to choose V? Adapt V according to arrival rate.

Token-bucket implemetation

$$\frac{dV}{dt} = \alpha(C - \lambda)$$

$$V(k+1) = [V(k) + tokens - packet\_size]^{+}$$

- No explicit measurements
- Code to implement this is a few lines

How to choose α?

- Fluid model for TCP: delay-differential equation.
- When the delay is zero, the roots of the characteristic equation are in the left-half plane?
- Find the largest value of the delay for which the roots hit the imaginary axis (root locus).

Design

$$\boxed{f(N,\alpha,d)=0}$$

- Fix N (number of users), d (RTT), gives an upper bound on α for local stability.
- Fix any two parameters, gives a bound on the free parameter.

II. Fair Multicast

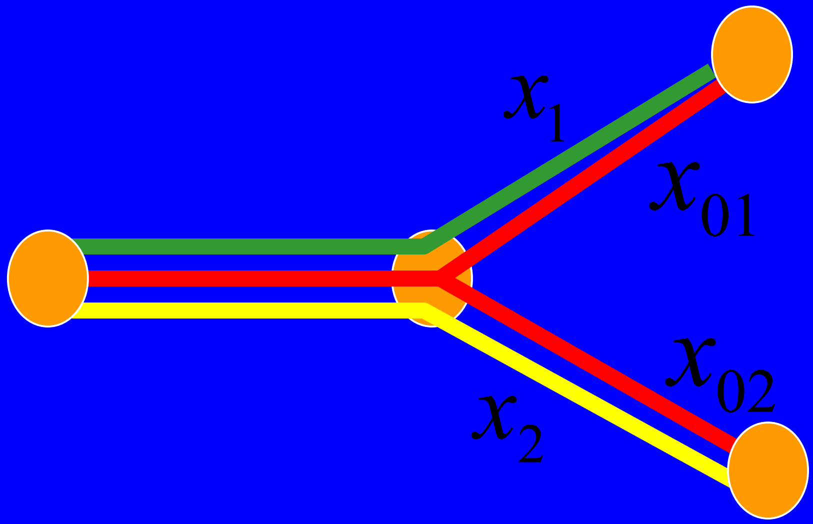

$$\max(x_{01}, x_{02}) + x_1 + x_2 \le C_A$$

(Fair multicast, continued)

Resource Allocation Problem

$$\max \quad U_{01}(x_{01}) + U_{02}(x_{02}) + U_1(x_1) + U_2(x_2)$$

$$\textit{subject to}$$

$$\max(x_{01}, x_{02}) + x_1 + x_2 \le C_A$$

$$x_{01} + x_1 \le C_B$$

$$x_{02} + x_2 \le C_C$$

$$\textit{All variables} \ge 0$$

Multicast Congestion Control

- TCP-type additive increase, multiplicative decrease algorithm
- Send marks only to those receivers that cause congestion at a link. One such receiver is randomly chosen.
- This algorithm converges to a penalty-function version of the convex program.

III. Pricing/Performance for Federated Networks

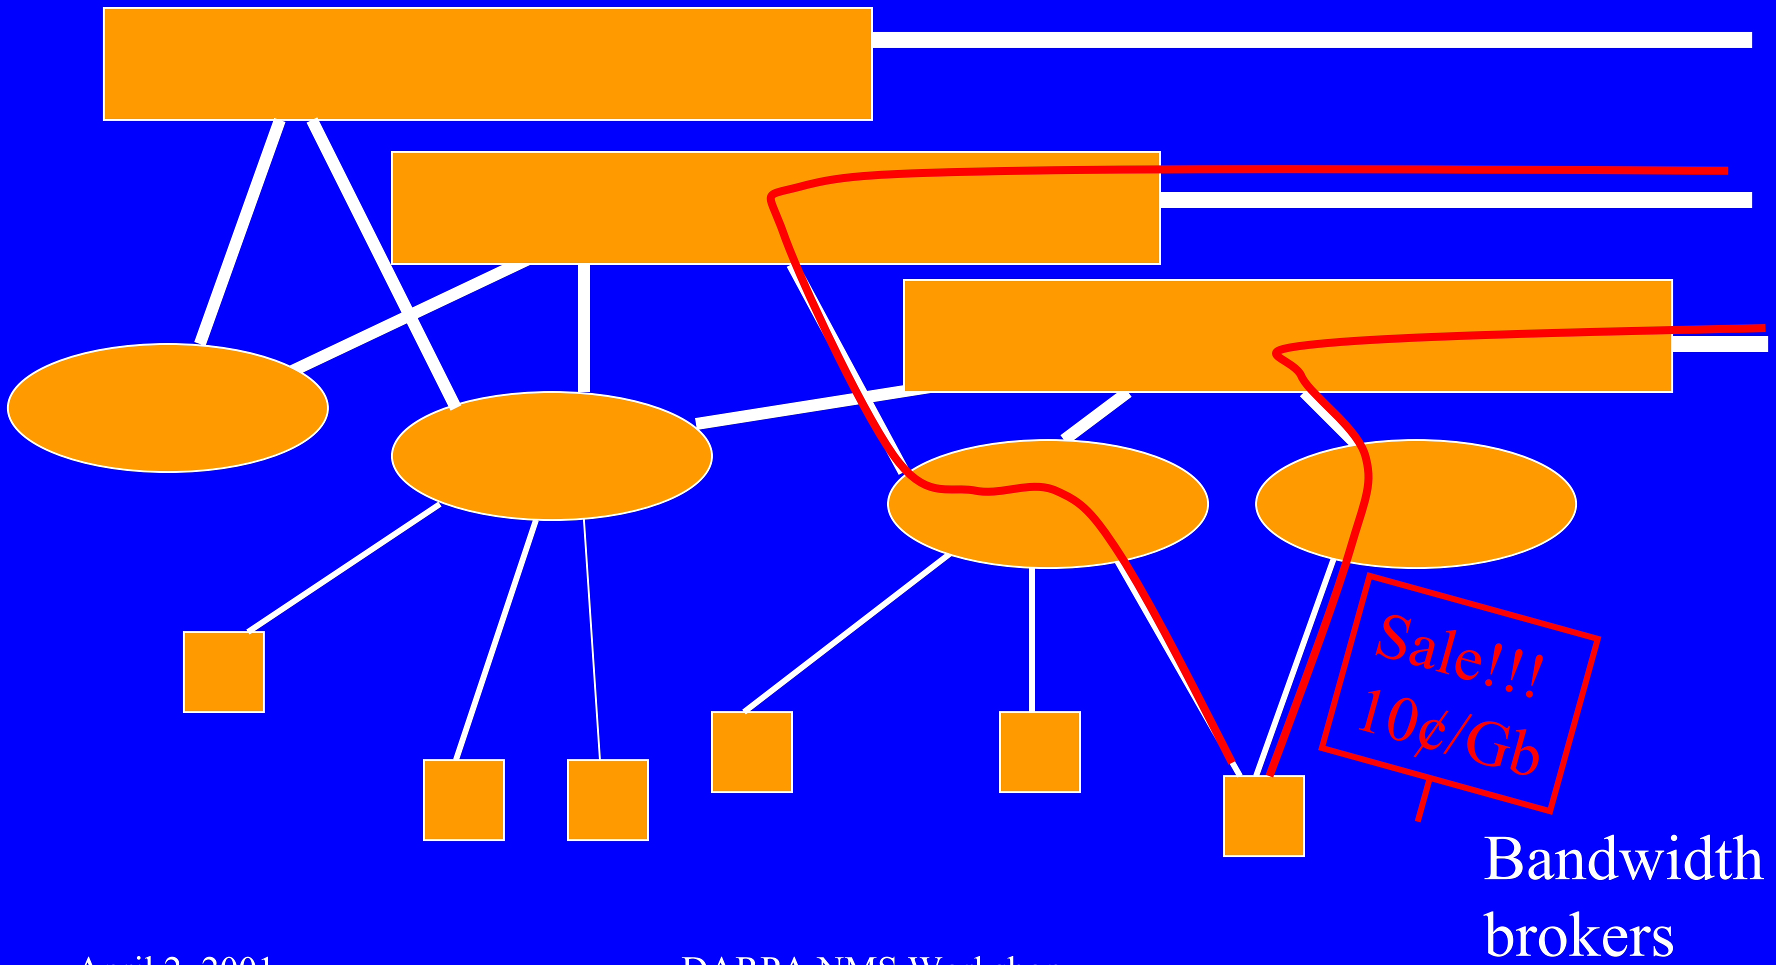

Some simple utility functions

Given U and price p, response is x that maximizes U(x)-xp.

Let $0<\theta<\infty$ and let $\alpha=(\theta-1)/\theta$

	Response function	Utility function	Expenditure
	$x(p)=p^{-\theta}$	$U(x)=x^{\alpha}/\alpha$	$p*x(p)=p^{1-\theta}$
$\theta<1$			
$\theta=1$			
$\theta>1$			

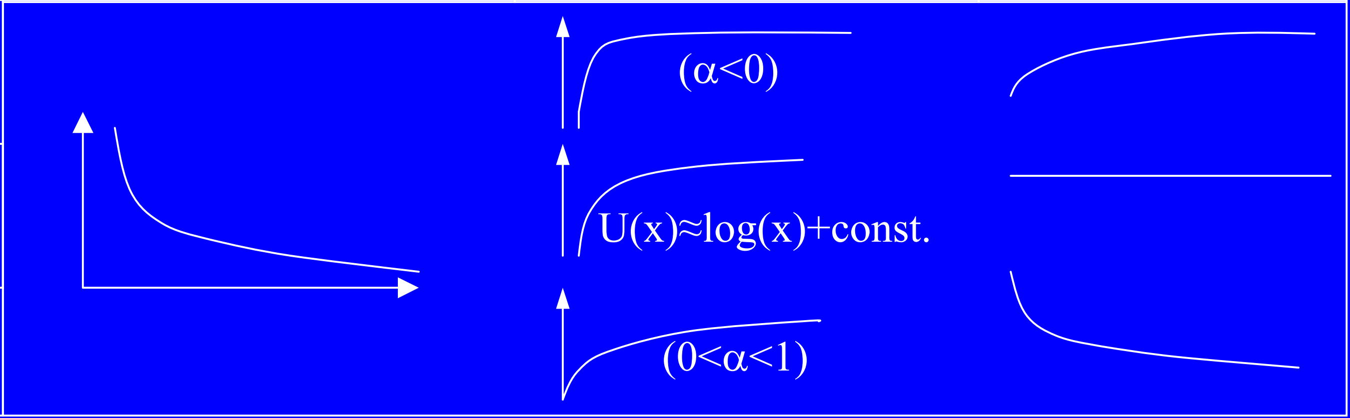

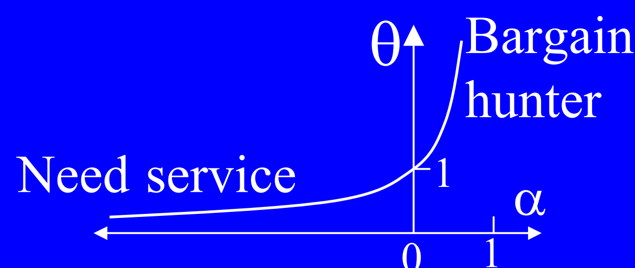

Price mark-up by middle agent

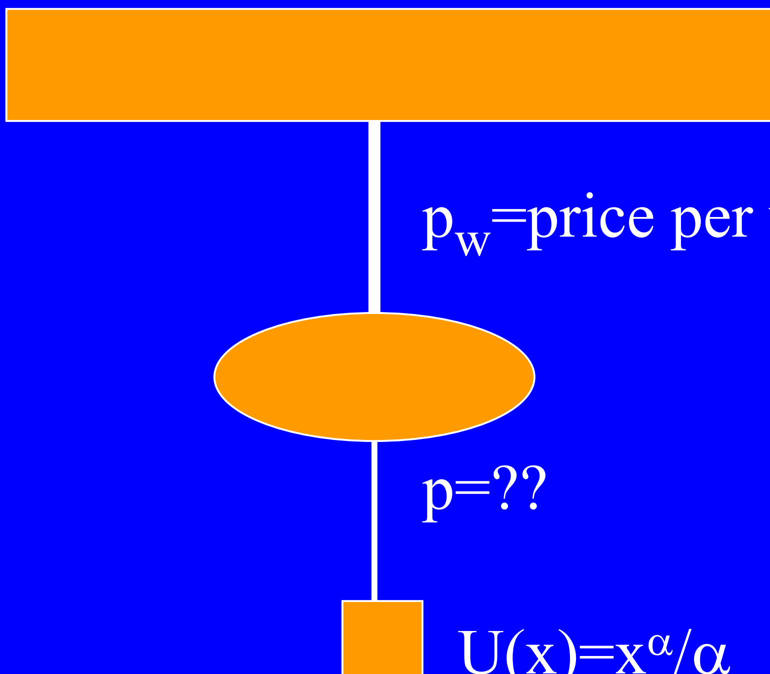

Q. What price p maximizes the profit $(p-p_w) * x(p)$ of middle agent/network?

A. $p=p_w/\alpha$ in case $0<\alpha<1$
$p=+\infty$ in case $\alpha \leq 0$

(Pricing in federated nets, continued)

Middle agents in series

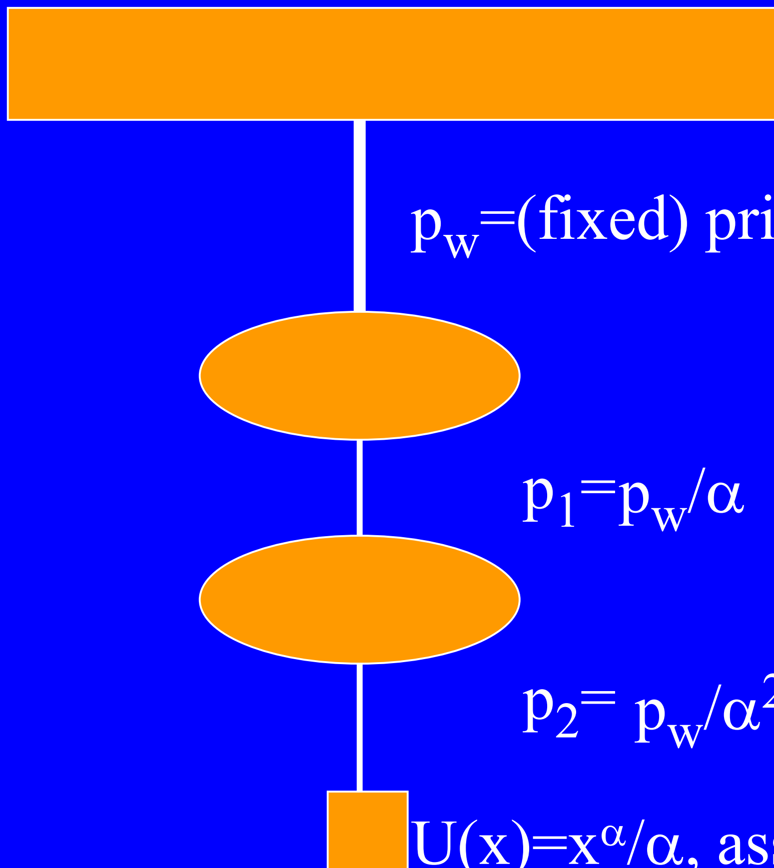

Each agent marks up the price by $1/\alpha$.
--Sum of profits is smaller than profit for a single middle agent.
--User pays higher price.
--Agent closer to user gets larger profit.

Fixed Path Networks

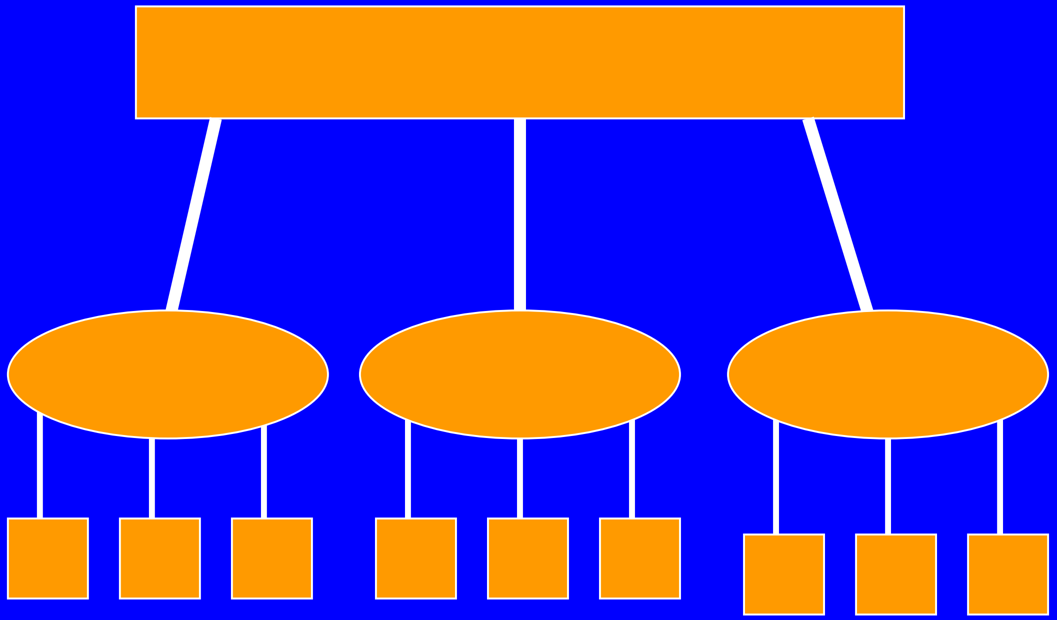

End users may have a mixture of utility functions.

Each network may have uniform pricing or targeted pricing, but overall effect is typically mixed unless per flow pricing is propagated through the network.

Middle agents in parallel

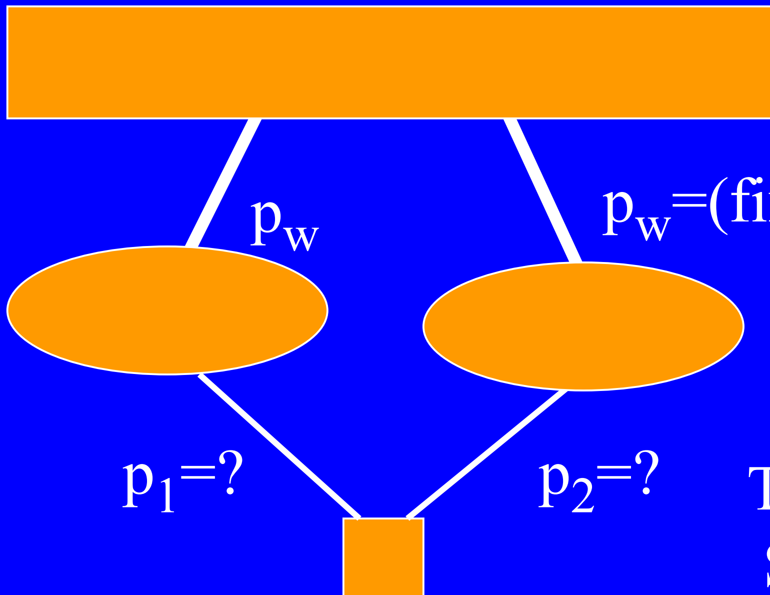

p_w=(fixed) price per unit traffic

There is no (Nash) equilibrium.
So what? Fixes?

$U(x)=x^\alpha/\alpha$, use price based routing

(i) If p_w's are even slightly different, then there is a Nash equilibrium.

(ii) Congestion -- links on a bottleneck are valuable--allocation of a finite resource kicks in. <u>Price to fill.</u>

(iii) Others: slow down oscillations through long-term contracts, delay in pricing information, etc.

Congestion based pricing

Single class, “core stateless”

Multiple classes (Diff serve)

Per flow QoS (ATM,IS)

Adopt “throughput is the only thing” philosophy. Allow bursty traffic by charging for packets causing congestion.

Provide low delay for premium charge rate.

Charge for QoS given input traffic constraints.

Paying Networks for Marks

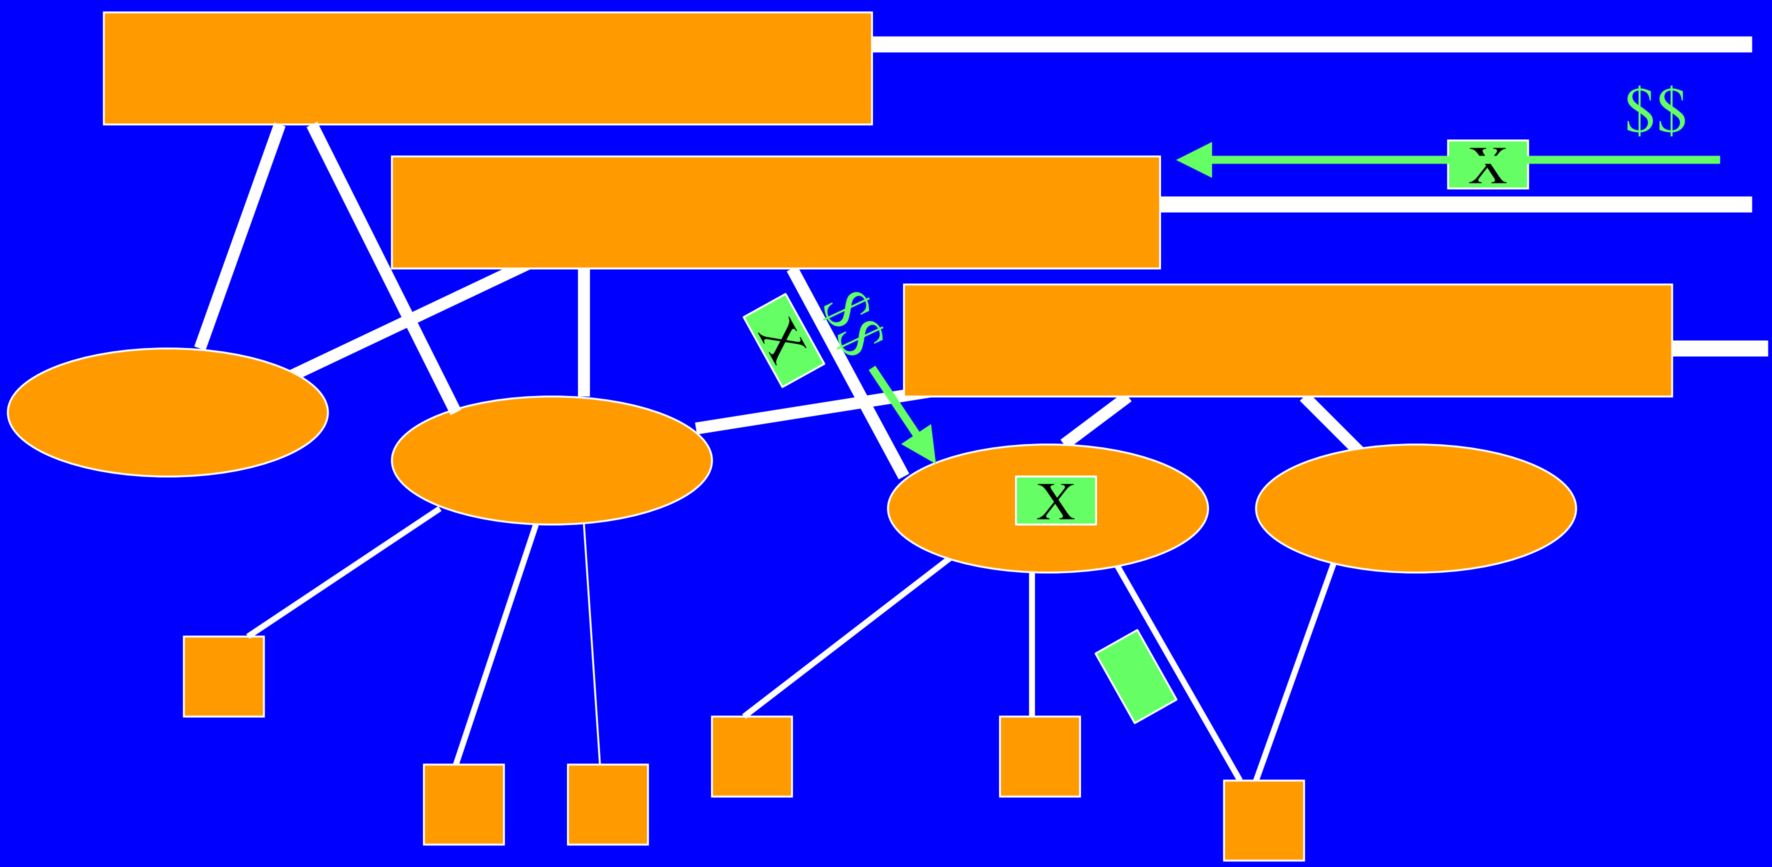

Parting thought on pricing

Q. Can the network utility be maximized if individual networks charge for congestion in order to control throughput to maximize their profits?

A. Perhaps.

-Competition for upper tier networks,

Keeps price near second lowest cost

Counters essentially unlimited bandwidth

Dampen oscillations by contracts

-Price to adjust demand to supply at bottlenecks

Summary

The work involves a hierarchy of network models.

- Stochastic Queueing Models at the Packet Level
- Fluid Models at the Flow Level
- Loss Models at the Connection Level
- Game-Theoretic Models of a Federation of Networks